

15/5/1 (Item 1 from file: 8)
DIALOG(R)File 8:Ei Compendex(R)
(c) 2003 Elsevier Eng. Info. Inc. All rts. reserv.

05760023 E.I. No: EIP01015487914

Title: **Methods for parallel execution of complex database queries**
Author: Reuter, Andreas
Corporate Source: Int Univ in Germany, Bruchsal, Ger
Source: Parallel Computing v 25 n 13 Dec 1999. p 2177-2188
Publication Year: 1999
CODEN: PACOEJ **ISSN:** 0167-8191
Language: English
Document Type: JA; (Journal Article) **Treatment:** A; (Applications); T; (Theoretical)

Journal Announcement: 0102W5

Abstract: During the last decade, all commercial database systems have included features for parallel processing into their products. This development has been driven by the fact that databases grow in size at considerable rates. According to the results of the 1998 'very large database contest' the world's largest databases, which have reached a size of over 10TB, double in size every year. At that speed, they outgrow the increase in processor speed and memory size, so additional measures are required to accommodate the effects of rapidly growing volumes of data. Parallelism is one of those options. It helps to keep processing times constant, even if the size of the **database increases**. That effect, which is often referred to as 'scaleup' is important for loading, index creation, all kinds of administrative operations on the database, and of course for long batch-type applications. Parallelism is also employed to speed-up queries that otherwise would take days or weeks to process and thus would be useless for the application. This type of requirement: fast results of complex queries on large data sets is characteristic of decision support applications. In this overview we will explain how parallelism in databases can help to solve such problems. (Author abstract) 19 Refs.

Descriptors: **Parallel** processing systems; Response time (computer systems); **Query** languages; Storage allocation (computer); Data structures ; Problem solving

Identifiers: Complex database queries

Classification Codes:

722.4 (Digital Computers & Systems); 723.3 (Database Systems); 722.1 (Data Storage, Equipment & Techniques); 723.2 (Data Processing)
722 (Computer Hardware); 723 (Computer Software)
72 (COMPUTERS & DATA PROCESSING)

15/5/4 (Item 4 from file: 8)
DIALOG(R)File 8:Ei Compendex(R)
(c) 2003 Elsevier Eng. Info. Inc. All rts. reserv.

04930509 E.I. No: EIP98024045884

Title: **On the application of parallel database technology for large scale document management systems**

Author: Clausnitzer, Alexander; Jaedicke, Michael; Mitschang, Bernhard; Nippl, Clara; Reiser, Angelika; Zimmermann, Stephan

Corporate Source: Technische Universitaet Muenchen, Munich, Ger

Conference Title: Proceedings of the 1997 International Database Engineering & Applications Symposium

Conference Location: Montreal, Can **Conference Date:** 19970825-19970827

Sponsor: Concordia University; IEEE

E.I. Conference No.: 47765

Source: Proceedings of the International Database Engineering & Applications Symposium, IDEAS 1997. IEEE, Piscataway, NJ, USA, 97TB100166. p 388-396

Publication Year: 1997

CODEN: 002754

Language: English

Document Type: CA; (Conference Article) **Treatment:** A; (Applications); G ; (General Review); T; (Theoretical)

Journal Announcement: 9804W1

Abstract: Modern document management systems administer an **increasing** amount of data. Unfortunately, traditional **database** technology cannot satisfy the resulting performance requirements. In this paper we want to report on how parallel database technology can contribute to make large and scalable document management systems work. Our discussion clearly shows that simply substituting a sequential database system by a parallel one does not necessarily lead to scalability and high performance. In contrast, detailed knowledge on the application and document management system is necessary to decide on the optimal architecture of the **parallel** database system, its **query** optimization and parallelization strategies, as well as on its data fragmentation scheme and load balancing properties. All discussions are exemplified referring to an existing document management system developed at our university that is currently migrated from a sequential database system to a parallel one. (Author abstract) 19 Refs.

Descriptors: Distributed database systems; **Parallel** processing systems; Management information systems; Computer architecture; **Query** languages; Optimization; Data reduction; Information retrieval systems

Identifiers: Parallel database technology; Large scale document management systems

Classification Codes:

723.3 (Database Systems); 722.4 (Digital Computers & Systems); 723.2 (Data Processing); 921.5 (Optimization Techniques)

723 (Computer Software); 722 (Computer Hardware); 921 (Applied Mathematics)

72 (COMPUTERS & DATA PROCESSING); 92 (ENGINEERING MATHEMATICS)

15/5/8 (Item 8 from file: 8)
DIALOG(R)File 8:Ei Compendex(R)
(c) 2003 Elsevier Eng. Info. Inc. All rts. reserv.

04247041 E.I. No: EIP95092850107

Title: Querying **genomic database** by using a parallel logic programming system on distributed computing environment

Author: Matsuda, Hideo

Corporate Source: Osaka Univ, Osaka, Jpn

Conference Title: Proceedings of the 1995 IEEE Pacific RIM Conference on Communications, Computers, and Signal Processing

Conference Location: Victoria, BC, Can Conference Date: 19950517-19950519

Sponsor: IEEE

E.I. Conference No.: 43519

Source: IEEE Pacific RIM Conference on Communications, Computers, and Signal Processing - Proceedings 1995. IEEE, Piscataway, NJ, USA, 95CH35765. p 333-336

Publication Year: 1995

CODEN: 002121

Language: English

Document Type: CA; (Conference Article) Treatment: A; (Applications); G; (General Review)

Journal Announcement: 9511W1

Abstract: We have developed a genomic database system on a network of workstations. The efficient access to genomic **databases** is crucial, given enormous **increase** in sequence data. By using a logic programming language, the system allows a user to perform adaptable data retrieval to integrated data objects in a single declarative framework. In addition by utilizing data-parallel processing, it provides efficient access in a large amount of genomic data on distributed computing environment. We present the implementation of the database system and its performance by measuring processing times for querying molecular sequence data. (Author abstract) 12 Refs.

Descriptors: **Parallel** processing systems; Logic programming; Distributed database systems; **Query** languages; Information **retrieval** ; Inference engines; Computer networks; Computer workstations; Performance

Identifiers: Genomic database; Parallel logic programming

Classification Codes:

723.4.1 (Expert Systems)

722.4 (Digital Computers & Systems); 723.1 (Computer Programming);

723.3 (Database Systems); 723.4 (Artificial Intelligence); 722.3 (Data Communication, Equipment & Techniques)

722 (Computer Hardware); 723 (Computer Software)

72 (COMPUTERS & DATA PROCESSING)

15/5/9 (Item 9 from file: 8)

DIALOG(R)File 8:EI Compendex(R)

(c) 2003 Elsevier Eng. Info. Inc. All rts. reserv.

04241495 E.I. No: EIP95092839922

Title: Pipelined disk arrays for digital movie retrieval

Author: Cohen, Ariel; Burkhard, Walter A.; Rangan, P. Venkat

Corporate Source: Univ of California, La Jolla, CA, USA

Conference Title: Proceedings of the International Conference on Multimedia Computing and Systems

Conference Location: Washington, DC, USA Conference Date: 19950515-19950518

Sponsor: IEEE

E.I. Conference No.: 43487

Source: International Conference on Multimedia Computing and Systems-Proceedings 1995. IEEE, Los Alamitos, CA, USA, 95TH8066. p 312-317

Publication Year: 1995

CODEN: 002114

Language: English

Document Type: CA; (Conference Article) Treatment: A; (Applications); T; (Theoretical)

Journal Announcement: 9510W5

Abstract: We develop a reliable disk array based storage architecture for digital video **retrieval**. Our goals are two-fold: maximizing the number of **concurrent** real-time sessions while minimizing the buffering requirements, and ensuring a high degree of reliability. The first goal is achieved by adopting a pipelined approach and by reducing latencies through specialized disk caching and constrained data placement schemes. The second goal is achieved by dividing the disks into RAID 3 reliability groups which serve as pipeline stages. We note that the buffering requirement decreases as the number of **groups increases**. To improve the performance further, we introduce two techniques for more efficient movie retrieval: on arrival caching, and interleaved annular layout. We present a case study of the performance of these techniques which shows a significant improvement when they are incorporated. (Author abstract) 9 Refs.

Descriptors: *Magnetic disk storage; Information retrieval; Pipeline processing systems; Digital image storage; Real time systems; Buffer storage; Performance; Bandwidth; ROM; Reliability

Identifiers: Pipelined disk arrays; Digital movie retrieval; Disk caching; Constrained data placement schemes

Classification Codes:

722.1 (Data Storage, Equipment & Techniques); 903.3 (Information Retrieval & Use); 722.4 (Digital Computers & Systems)

722 (Computer Hardware); 903 (Information Science)

72 (COMPUTERS & DATA PROCESSING); 90 (GENERAL ENGINEERING)

15/5/32 (Item 9 from file: 2)

DIALOG(R)File 2:INSPEC

(c) 2003 Institution of Electrical Engineers. All rts. reserv.

4589744 INSPEC Abstract Number: C9403-6160-016

Title: Parallel query processing

Author(s): Yu, P.S.; Chen, M.-S.; Wolf, J.L.; Turek, J.

Author Affiliation: Div. of Res., IBM Thomas J. Watson Res. Center, Yorktown Heights, NY, USA

Book Title: Advanced database systems p.229-58

Editor(s): Adam, N.R.; Bhargava, B.K.

Publisher: Springer-Verlag, Berlin, Germany

Publication Date: 1993 Country of Publication: West Germany xv+451

pp.

ISBN: 3 540 57507 3

Language: English Document Type: Book Chapter (BC)

Treatment: Practical (P)

Abstract: With the advent of inexpensive microprocessors and high bandwidth interconnects, coupling a large number of processors to form a highly parallel system has become an increasingly popular method for improving the cost-performance ratio of computer systems. Recent work has shown that this method is also applicable to **database** systems with **increasing** benefits as the queries become larger and more complex. The objective of this paper is to examine the various issues encountered in **parallel query** processing as well as the techniques that are available for addressing these issues. (0 Refs)

Subfile: C

Descriptors: database management systems; parallel processing; query processing

Identifiers: microprocessors; high bandwidth interconnects; cost-performance ratio; **parallel query** processing; database

Class Codes: C6160 (Database management systems (DBMS)); C6150N (Distributed systems); C5440 (Multiprocessor systems and techniques)

15/5/34 (Item 11 from file: 2)

DIALOG(R)File 2:INSPEC

(c) 2003 Institution of Electrical Engineers. All rts. reserv.

03960762 INSPEC Abstract Number: C91058164

Title: High speed text retrieval from large databases on a massively parallel processor

Author(s): Reddaway, S.F.

Author Affiliation: Active Memory Technol. Ltd., Reading, UK

Journal: Information Processing & Management vol.27, no.4 p.311-16

Publication Date: 1991 Country of Publication: UK

CODEN: IPMADK ISSN: 0306-4573

U.S. Copyright Clearance Center Code: 0306-4573/91/\$3.00+.00

Language: English Document Type: Journal Paper (JP)

Treatment: Practical (P)

Abstract: A word-based text document retrieval system is being implemented by AMT (Active Memory Technology) on the AMT DAP series of massively parallel computers. For a 20 G Byte database of 10 million documents a system with 2 DAP 610s can sustain around 70 Boolean queries per second. Any of these parameters can be extended. In the initial application the addition of two DAP 610s to a **cluster** of four large VAXes will **increase** the performance by several hundred times. (10 Refs)

Subfile: C

Descriptors: information retrieval; parallel architectures

Identifiers: DAP 610; VAX; text document retrieval system; AMT; Active Memory Technology; AMT DAP; parallel computers; database; Boolean queries; 20 Gbytes

Class Codes: C5440 (Multiprocessor systems and techniques); C7250 (Information storage and retrieval)

Numerical Indexing: memory size 2.1E+10 Byte

15/5/36 (Item 2 from file: 233)

DIALOG(R)File 233:Internet & Personal Comp. Abs.

(c) 2003, EBSCO Pub. All rts. reserv.

00496307 98BY05-013

When parallel lines meet -- Only parallel databases bring the best performance out of parallel-processing systems

Rudin, Ken

BYTE, May 1, 1998, v23 n5 p81-88, 7 Page(s)

ISSN: 0360-5280

Languages: English

Document Type: Articles, News & Columns

Geographic Location: United States

Discusses how parallel databases bring out the best performance from parallel processing systems. Illustrates how partitioning and pipelining **increase database** performance in a sample sequential query. Says that

types of multiprocessor architectures include shared-everything, shared-disk and shared-nothing architectures. Says that there are also two kinds of parallel database architectures, shared-nothing and shared-disk, which were intended to address issues in shared-nothing hardware designs. Concludes that the use of **parallel queries** offer the most effective approach to improve **query** -processing performance. Includes a sidebar "Partitioning Data" about the basic methods of partitioning data. Includes three diagrams, two graphs and one flow chart.

Descriptors: Database; Parallel Processing; Data Base Management; Architecture

15/5/40 (Item 6 from file: 233)

DIALOG(R)File 233:Internet & Personal Comp. Abs.

(c) 2003, EBSCO Pub. All rts. reserv.

00348620 94IW05-104

Oracle set to make venture into parallel universe with Version 7.1

Mace, Scott

InfoWorld , May 9, 1994 , v16 n19 p1, 118, 2 Page(s)

ISSN: 0199-6649

Company Name: Oracle

Product Name: Oracle

Languages: English

Document Type: Product Announcement

Geographic Location: United States

Reports that Oracle Corp. of Redwood Shores, CA, will release Oracle 7.1 (\$NA). Says it will offer user-extensible SQL with application-specific data types and function, a primary site replication, and a GUI-based server management product. Says a **Parallel Query** Option will also be available for Sun Microsystems Inc. and Sequent Computer Systems Inc. symmetric multiprocessing servers, and will ship on 12 additional parallel processing s soon. **Adds** that Oracle's synchronous **replication** feature, whi allows multiple servers to gather and replicate data, has entered beta testing. (cr)

Descriptors: Data Base Management; **Parallel** Processing; Server; Structured **Query** Language; Database; Multiprocessing
Identifiers: Oracle; Oracle

15/5/48 (Item 1 from file: 95)

DIALOG(R)File 95:TEME-Technology & Management

(c) 2003 FIZ TECHNIK. All rts. reserv.

00562826 I92013435927

Dynamic concurrency control algorithms for large distributed database systems

(Dynamische Algorithmen der Zugriffssteuerung fuer grosse verteilte Datenbanksysteme)

Ghafoor, A; Farhat, FY

Sch. of Electr. Eng., Purdue Univ., West Lafayette, IN, USA

Computer Journal, London, v34, n6, ppA095-A104, 1991

Document type: journal article Language: English

Record type: Abstract

ISSN: 0010-4620

ABSTRACT:

The authors propose two concurrency control algorithms for distributed database systems. Both algorithms employ a token passing mechanism, in which the path of the token is determined dynamically in order to maximize the system's throughput. The two algorithms differ in terms of selecting the next holder of the token in case the system becomes empty. One scheme, the Dynamic-Ring Algorithm (DRA) switches to a virtual ring algorithm. The other, the Dynamic-Status Algorithm (DSA) employs a query based mechanism. Simulation study reveals that both the proposed algorithms perform considerably better than the conventional ring algorithm, under low to medium loading conditions in the network. The difference in performance becomes more prominent as the number of **database sites increases** .

Summary only published. Full paper available from journal editor.

DESCRIPTORS: ACCESS CONTROL; ALGORITHM; **QUERY** LANGUAGES; ALGORITHM THEORY
; DISTRIBUTED DATABASES; **CONCURRENCY** CONTROL
IDENTIFIERS: LARGE DISTRIBUTED DATABASE SYSTEMS; CONCURRENCY CONTROL
ALGORITHMS; TOKEN PASSING MECHANISM; DYNAMIC RING ALGORITHM; VIRTUAL RING
ALGORITHM; DYNAMIC STATUS ALGORITHM; QUERY BASED MECHANISM;
Zugriffssteuerung; verteiltes Datenbanksystem

15/5/49 (Item 2 from file: 95)
DIALOG(R)File 95:TEME-Technology & Management
(c) 2003 FIZ TECHNIK. All rts. reserv.

00547391 E92024144080

Towards parallelism in a relational database system
(Ueber Parallelitaeten in einem relationalen Datenbanksystem)
Loibl, E; Obermaier, H; Pawlowski, M
TU Muenchen, D
1991
Document type: Report Language: English
Record type: Abstract

ABSTRACT:

To fit the needs of nowadays complex **database** applications it is necessary to **increase** the performance of relational **database** systems dramatically. This challenge can only be achieved by use of massive parallel systems. In this paper we want to show an methodical way of parallelizing an existing database system for a general-purpose multiprocessor. Four approaches of different granularity, i.e. kernel replication, module splitting, query processing, and operator parallelism, are discussed and their application to the database system TransBase is shown.